

(captioned "Version With Markings To Show Changes Made To Claims Relative To U.S. Patent No. 5,857,180") in accordance with 37 CFR § 1.173(g) and a marked-up version of the claims that indicates the differences between the current set of claims and the state of the claims prior to this amendment (captioned "Version With Markings To Show Changes Made To Claims in This Amendment") for ease of examination. A discussion of the status of the pending claims as required by 37 CFR § 1.173(c) is provided below. The following is a clean version of the claims amended herein:

1 27. (Twice Amended) A method of parallelizing an operation, the method comprising the
2 steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions; wherein the step of assigning
7 work partitions is performed by assigning the work partitions in a sequence based
8 at least in part on sizes associated with the work partitions; and
9 said plurality of entities operating in parallel on work partitions assigned to them to
10 perform said operation.

1 29. (Twice Amended) A method of parallelizing an operation, the method comprising the
2 steps of:
3 dividing the operation into a set of work partitions;

4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions, wherein the step of assigning
7 work partitions includes:
8 assigning said at least one entity a first work partition from said set of work
9 partitions; and
10 after said at least one entity has completed operation on said first work partition,
11 assigning said at least one entity a second work partition from said set of
12 work partitions, wherein the step of assigning said at least one entity a
13 second work partition includes
14 determining whether there are any unassigned work partitions from a first
15 level in the hierarchy to which said first work partition belonged;
16 and
17 if there are no unassigned work partitions from the first level in the
18 hierarchy, then selecting said second work partition from a level in
19 said hierarchy that is two levels above said first level in said
20 hierarchy;
21 said plurality of entities operating in parallel on work partitions assigned to them to
22 perform said operation; and
23 wherein the operation is specified in a query that corresponds to a hierarchy of
24 operations.

1 30. (Once Amended) A method of parallelizing an operation, the method comprising the
2 steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation;
9 the method includes the step of generating a serial execution plan for operations in a
10 database management system (DBMS) running on a computer system;
11 the method includes the step of generating a parallelized execution plan for said serial
12 execution plan, said parallelized execution plan including first and second
13 operations;
14 the step of dividing an operation is performed by dividing said second operation;
15 the plurality of entities includes one or more slave processes operating on a plurality of
16 data partitions, the quantity of said data partitions being greater than the quantity
17 of said slave processes;
18 executing said parallelized execution plan when a plurality of parallel resources of said
19 computer system are available; and
20 executing said serial execution plan when said plurality of resources are not available.

1 31. (Once amended) The method of claim 30 wherein said step of generating a parallelized
2 execution plan includes the steps of:

identifying one or more segments of said serial execution plan that can be parallelized;
and
identifying partitioning requirements of said one or more segments.

33. (Once Amended) A method of parallelizing an operation, the method comprising the
steps of:
dividing the operation into a set of work partitions;
assigning work partitions from said set of work partitions to a plurality of entities,
wherein at least one entity of said plurality of entities is assigned a plurality of
work partitions from said set of work partitions;
said plurality of entities operating in parallel on work partitions assigned to them to
perform said operation;
generating an execution plan for said operation;
examining said execution plan from bottom up;
identifying a parallelized portion of said execution plan, said parallelized portion can be
processed in parallel, said parallelized portion including first and second
operations, said first and second operations being executable in parallel;
wherein the step of dividing the operation is performed by dividing said second
operation;
wherein the plurality of entities includes one or more slave processes operating on a
plurality of data partitions, the quantity of said data partitions being greater than
the quantity of said slave processes;

19 identifying some serial portion of said execution plan, said serial portion can be
20 processed in serial; and
21 allocating a central scheduler between said parallelized portion and said serial portion.

1 35. (Once amended) A method for parallelizing an operation, the method comprising the
2 steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation;
9 generating an execution plan to execute database management system (DBMS)
10 operations in parallel, said execution plan including first and second operations;
11 wherein the step of dividing said operation is performed by dividing said second
12 operation;
13 initiating an operation coordinator in a computer system to coordinate execution of said
14 execution plan;
15 initiating, by said operation coordinator, a first set of slaves operating on a plurality of
16 data partitions to produce data, the quantity of said data partitions being greater
17 than the quantity of said first set of slave processes;
18 initiating, as said plurality of entities, by said operation coordinator, a second set of
19 slaves to consume data; and

20 directing said second set of slaves to produce data and said first set of slaves to consume
21 data when said first set of slaves finishes producing data.

1 36. (Once amended) The method of claim 35 wherein said execution plan is comprised of
2 operator nodes and said operator nodes are linked together to form execution sets.

1 37. (Once amended) A method for parallelizing an operation, the method comprising the
2 steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation;
9 generating an execution plan to execute said operations in parallel, said execution plan
10 including first and second operations;
11 wherein the step of dividing said operation includes dividing said first operation;
12 initiating producer slaves operating on a plurality of data partitions to produce a first data
13 production;
14 initiating consumer slaves to consume said first data production;
15 when said first data production is completed, generating an identification of a plurality of
16 said consumer slaves that did not receive data in said first data production;
17 examining said identification during a subsequent data production; and

18 reducing said subsequent data production such that said subsequent data production does
19 not produce data for said plurality of said consumer slaves.

D⁵
1 44. (Once amended) The method of Claim 43 wherein the step of incorporating hints that
2 dictate the operation of a table scan includes incorporating hints that rowed partitioning is
3 to be used during the table scan.

D⁶
1 63. (Once amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions; wherein the step of assigning
7 work partitions is performed by assigning the work partitions in a sequence based
8 at least in part on sizes associated with the work partitions; and
9 said plurality of entities operating in parallel on work partitions assigned to them to
10 perform said operation.

D⁷
1 65. (Twice Amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of

6 work partitions from said set of work partitions, wherein the step of assigning
7 work partitions includes
8 assigning said at least one entity a first work partition from said set of work
9 partitions; and
10 after said at least one entity has completed operating on said first work partition,
11 assigning said at least one entity a second work partition from said set of
12 work partitions;
D 13 said plurality of entities operating in parallel on work partitions assigned to them to
14 perform said operation;
15 wherein the operation is specified in a query that corresponds to a hierarchy of
16 operations; and
17 the step of assigning said at least one entity a second work partition includes
18 determining whether there are any unassigned work partitions from a first level in
19 the hierarchy to which said first work partition belonged; and
20 if there are no unassigned work partitions from the first level in the hierarchy,
21 then selecting said second work partition from a level in said hierarchy
22 that is two levels above said first level in said hierarchy.

1 66. (Once amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;

4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operation in parallel on work partitions assigned to them to
8 perform said operation;
9 wherein the instructions include instructions for performing the step of generating a serial
10 execution plan for operations in a database management system (DBMS) running
11 on a computer system;
12 wherein the instructions include instructions for performing the step of generating a
13 parallelized execution plan for said serial execution plan, said parallelized
14 execution plan including first and second operations;
15 wherein the step of dividing an operation is performed by dividing said second operation;
16 wherein the plurality of entities includes one or more slave processes operating on a
17 plurality of data partitions, the quantity of said data partitions being greater than
18 the quantity of said slave processes;
19 wherein the instructions include instructions for performing the step of executing said
20 parallelized execution plan when a plurality of parallel resources of said computer
21 system are available; and
22 wherein the instructions include instructions for performing the step of executing said
23 serial execution plan when said plurality of resources are not available.

1 67. (Once amended) The computer-readable medium of claim 66 wherein said step of
2 generating a parallelized execution plan includes the steps of:

3 identifying one or more segments of said serial execution plan that can be parallelized;
4 and
5 identifying partitioning requirements of said one or more segments.

D¹
1 68. (Once amended) The computer-readable medium of claim 66 wherein said step of
2 generating a parallelized execution plan is based on a specification of parallelism in a
3 statement specifying one of said operations.

1 69. (Once amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform some operation;
9 generating an execution plan for said operation;
10 examining said execution plan from bottom up;
11 identifying a parallelized portion of said execution plan, said parallelized portion can be
12 processed in parallel, said parallelized portion including first and second
13 operations, said first and second operations being executable in parallel;
14 wherein the step of dividing the operation is performed by dividing said second
15 operation;

16 wherein the plurality of entities includes one or more slave processes operating on a
17 plurality of data partitions, the quantity of said data partitions being greater than
18 the quantity of said slave processes;
19 identifying some serial portion of said execution plan, said serial portion can be
20 processed in serial; and
21 allocating a central scheduler between said parallelized portion and said serial portion.

51 1 70. (Once amended) The computer-readable medium of Claim 69 further including
2 instructions for performing the steps of:
3 identifying a first data flow requirement for a first portion of said execution plan said first
4 data flow requirement corresponding to a partitioning of a data flow required by
5 said first portion;
6 identifying a second data flow requirement for a second portion of said execution plan
7 said second data flow requirement corresponding by said second portion; and
8 allocating a data flow director between said first portion and said second portion when
9 said first data flow requirement is not compatible with said second data flow
10 requirement said data flow director repartitioning a data flow of said first portion
11 to be compatible with said second data flow requirement.

1 71. (Once amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;

4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation;
9 generating an execution plan to execute database management system (DBMS)
10 operations in parallel, said execution plan including first and second operations;
D¹ 11 wherein the step of dividing said operation is performed by dividing said second
12 operation;
13 initiating an operation coordinator in a computer system to coordinate execution of said
14 execution plan;
15 initiating, by said operation coordinator, a first set of slaves operating on a plurality of
16 data partitions to produce data, the quantity of said data partitions being greater
17 than the quantity of said first set of slave processes;
18 initiating, as said plurality of entities, by said operation coordinator, a second set of
19 slaves to consume data; and
20 directing said second set of slaves to produce data and said first set of slaves to consume
21 data when said first set of slaves finishes producing data.

1 72. (Once amended) The computer-readable medium of claim 71 wherein said execution plan
2 is comprised of operator nodes and said operator nodes are linked together to form
3 execution sets.

D⁷

1 73. (Once amended) A computer-readable medium carrying instructions for parallelizing an
2 operation, the instructions including instructions for performing the steps of:
3 dividing the operation into a set of work partitions;
4 assigning work partitions from said set of work partitions to a plurality of entities,
5 wherein at least one entity of said plurality of entities is assigned a plurality of
6 work partitions from said set of work partitions;
7 said plurality of entities operating in parallel on work partitions assigned to them to
8 perform said operation;
9 generating an execution plan to execute said operations in parallel, said execution plan
10 including first and second operations;
11 wherein the step of dividing said operation includes dividing said first operation;
12 initiating producer slaves operating on a plurality of data partitions to produce a first data
13 production;
14 initiating consumer slaves to consume said first data production;
15 when said first data production is completed, generating an identification of a plurality of
16 said consumer slaves that did not receive data in said first data production;
17 examining said identification during a subsequent data production; and
18 reducing said subsequent data production such that said subsequent data production does
19 not produce data for said plurality of said consumer slaves.

D⁸

1 80. (Once amended) The computer-readable medium of Claim 79 wherein the step of
2 incorporating hints that dictate the operation of a table scan includes incorporating hints
3 that rowed partitioning is to be used during the table scan.
